

CS422 - Programming Language Semantics

Lambda Calculus and Combinatory Logic

Grigore Roşu

Department of Computer Science
University of Illinois at Urbana-Champaign

Here we discuss the relationships between two important and closely related mathematical theories:

- *Lambda calculus*, written also *λ -calculus*, is a pure calculus of functional abstraction and function application, with many applications in logic and computer science;
- *Combinatory logic* shows that bound variables can be entirely eliminated without loss of expressiveness. It has applications both in the foundations of mathematics and in the implementation of functional programming languages.

A good reference for these subjects is the book “*The Lambda Calculus: Its Syntax and Semantics*” by H.P. Barendregt (*Second Edition, North Holland 1984*). This book also contains a great discussion on the history and motivations of these theories.

Lambda Calculus (λ -Calculus)

Lambda calculus was introduced in 1930s, as a mathematical theory together with a proof calculus aiming at capturing foundationally the important notions of function and function application. Those years were marked by several paradoxes in mathematics and logics. The original motivation of λ -calculus was to provide a foundation of logics and mathematics. Even though the issue of whether λ -calculus indeed provides a strong foundation of mathematics is still largely open, it nevertheless turned out to be a quite successful *theory of computation*.

Syntax of λ -Calculus

Assume an infinite set of *variables*, or *names*, V . Then the syntax of *λ -expressions* is (in BNF notation)

$$Exp ::= Var \mid Exp \ Exp \mid \lambda Var. Exp$$

where Var ranges over the variables in V . We will use lower letters x, y, z , etc., to refer to variables, and capital letters E, E', E_1, E_2 , etc., to refer to λ -expressions. The following are therefore examples of λ -expressions: $\lambda x.x$, $\lambda x.xx$, $\lambda x.(fx)(gx)$, $(\lambda x.fx)x$.

λ -Expressions of the form $\lambda x.E$ are called *λ -abstractions*, and those of the form $E_1 E_2$ are called *λ -applications*. The former captures the intuition of “function definition”, while the latter that of “function application”. To avoid parentheses, assume that λ -application is left associative and binds tighter than λ -abstraction.

Many programming language concepts, and even entire programming languages, translate relatively naturally into λ -calculus concepts or into λ -expressions. In particular, one can define some transformation φ from functional expressions into λ -expressions. Such a transformation φ would take, for example,

- variable names x to unique variables $x \in Var$;
- function declarations of the form `fun x -> E` to $\lambda x.\varphi(E)$; and
- bindings (which generalize the idea of “local declarations” occurring in most programming languages, functional or not) `let x1 = E1 and x2 = E2 and ... and xn = En in E` to λ -expressions $(\lambda x_1.\lambda x_2.\cdots.\lambda x_n.\varphi(E))\varphi(E1)\varphi(E2)\cdots\varphi(En)$.

Free and Bound Variables

Variable occurrences in λ -expressions can be either *free* or *bound*. Given a λ -abstraction $\lambda x.E$, also called a *binding*, then the variable x is said to be *declared* by the λ -abstraction, or that λ *binds x in E* ; also, E is called the *scope* of the binding.

Formally, we define the set $FV(E)$ of *free variables of E* as follows:

- $FV(x) = \{x\}$,
- $FV(E_1 E_2) = FV(E_1) \cup FV(E_2)$, and
- $FV(\lambda x.E) = FV(E) - \{x\}$.

Consider the three underlined occurrences of x in the λ -expression $(\lambda \underline{x}.\lambda y.y \underline{x} y)\underline{x}$. The first is called a *binding occurrence of x* , the second a *bound occurrence of x* (this occurrence of x is bound to the binding occurrence of x), and the third a *free occurrence of x* . Expressions E with $FV(E) = \emptyset$ are called *closed* or *combinators*.

Substitution

Evaluation of λ -expressions is “by-substitution”. That means that the λ -expression that is “passed” to a λ -abstraction is “copied as it is” at all the bound occurrences of the binding variable. This will be formally defined later. Let us now formalize and discuss the important notion of *substitution*. Intuitively, $E[x \leftarrow E']$ represents the λ -expression obtained from E by replacing each free occurrence of x by E' . Formally, substitution can be defined as follows:

- $y[x \leftarrow E'] = \begin{cases} E' & \text{if } y = x \\ y & \text{if } y \neq x, \end{cases}$
- $(E_1 E_2)[x \leftarrow E'] = (E_1[x \leftarrow E'])(E_2[x \leftarrow E']),$
- $(\lambda x.E)[x \leftarrow E'] = \lambda x.E.$

The tricky part is to define substitution on λ -abstractions of the

form $(\lambda y.E)[x \leftarrow E']$, where $y \neq x$. That is because E' may contain free occurrences of y ; these occurrences of y would become bound by the binding variable y if one simply defined this substitution as $\lambda y.(E[x \leftarrow E'])$ (and if E had any free occurrences of x), thus violating the intuitive meaning of binding. This phenomenon is called *variable capturing*. Consider, for example, the substitution $(\lambda y.x)[x \leftarrow yy]$; if one applies the substitution blindly then one gets $\lambda y.yy$, which is most likely *not* what one meant (since $\lambda y.x$ is by all means “equivalent” to $\lambda z.x$ - this equivalence will be formalized shortly - while $\lambda y.yy$ is not equivalent to $\lambda z.yy$). There are at least three approaches in the literature to deal with this delicate issue:

1. Define $(\lambda y.E)[x \leftarrow E']$ as $\lambda y.(E[x \leftarrow E'])$, but pay special attention whenever substitution is used to add sufficient conditions to assure that y is not free in E' . This approach simplifies the definition of substitution, but complicates the presentation of λ -calculus by having to mention “obvious”

additional hypotheses all the time a substitution is invoked;

2. Define substitution as a *partial operation*: $(\lambda y.E)[x \leftarrow E']$ is defined and equal to $\lambda y.(E[x \leftarrow E'])$ if and only if $y \notin FV(E')$. This may seem like the right approach, but unfortunately is also problematic, because the entire equational definition of λ -calculus would then become *partial*, which has serious technical implications w.r.t. mechanizing equational deduction (or the process of proving λ -expressions equivalent) and rewriting (or reduction);
3. Define substitution as a *total operation*, but apply a renaming of y to some variable that does not occur in E or E' in case $y \in FV(E')$ (this renaming is called *α -conversion* and will be defined formally shortly). This approach slightly complicates the definition of substitution, but simplifies the presentation of many results later on. It is also useful when one wants to mechanize λ -calculus, because it provides an algorithmic way

to avoid variable capturing:

$$(\lambda y.E)[x \leftarrow E'] = \begin{cases} \lambda y.(E[x \leftarrow E']) & \text{if } y \notin FV(E') \\ \lambda z.((E[y \leftarrow z])[x \leftarrow E']) & \text{if } y \in FV(E') \end{cases},$$

where z is a new variable that does not occur in E or E' . Note that the requirement “ z does not occur in E or E' ” is stronger than necessary, but easier to state that way.

All three approaches above have their advantages and disadvantages, and one can find many scientists defending each of them. However, we will later on choose a totally different approach to define λ -calculus as an executable specification, in which substitutions play no role anymore. More precisely, we will define λ -calculus through its translation to combinatory logic.

α-Conversion

In mathematics, functions that differ only in the name of their variables are equal. For example, the functions f and g defined (on the same domain) as $f(x) = x$ and $g(y) = y$ are considered identical. However, with the machinery developed so far, there is no way to show that the λ -expressions $\lambda x.x$ and $\lambda y.y$ are equal. It is a common phenomenon in the development of mathematical theories to add desirable but unprovable properties as axioms. The following is the first meaningful equational axiom of λ -calculus, known under the name of *α-conversion*:

$$(\alpha) \quad \lambda x.E = \lambda z.(E[x \leftarrow z])$$

for any variable z that does not occur in E (this requirement on z is again stronger than necessary, but it is easier to state).

Using the equation above, one has now the possibility to prove λ -expressions “equivalent”. To capture this provability relation formally, we let $E \equiv_\alpha E'$ denote the fact that the equation $E = E'$ can be proved using standard equational deduction from the equational axioms above ((α) plus those for substitution).

Exercise 1 *Prove the following equivalences of λ -expressions:*

- $\lambda x.x \equiv_\alpha \lambda y.y$,
- $\lambda x.x(\lambda y.y) \equiv_\alpha \lambda y.y(\lambda x.x)$,
- $\lambda x.x(\lambda y.y) \equiv_\alpha \lambda y.y(\lambda y.y)$.

β-Equivalence and β-Reduction

We now define another important equation of λ -calculus, known under the name of *β-equivalence*:

$$(\beta) \quad (\lambda x.E)E' = E[x \leftarrow E']$$

The equation (β) tells us how λ -abstractions are “applied”. Essentially, it says that the argument λ -expression that is passed to a λ -abstraction is copied at every free occurrence of the variable bound by the λ -abstraction within its scope.

We let $E \equiv_\beta E'$ denote the fact that the equation $E = E'$ can be proved using standard equational deduction from the equational axioms above: (α) , (β) , plus those for substitution. For example $(\lambda f.fx)(\lambda y.y) \equiv_\beta x$, because one can first deduce that $(\lambda f.fx)(\lambda y.y) = (\lambda y.y)x$ by (β) and then that $(\lambda y.y)x = x$ also by

(β) ; the rest follows by the transitivity rule of equational deduction.

Exercise 2 Show that $(\lambda x.(\lambda y.x))yx \equiv_{\beta} y$

When the equation (β) is applied only from left to right, that is, as a rewrite rule, it is called *β -reduction*. We let \Rightarrow_{β} denote the corresponding rewriting relation on λ -expressions. To be more precise, the relation \Rightarrow_{β} is defined on α -equivalence classes of λ -expressions; in other words, \Rightarrow_{β} applies *modulo α -equivalence*.

Given a λ -expression E , one can always apply α -conversion on E to rename its binding variables so that all these variables have different names which do not occur in $FV(E)$. If that is the case, then note that variable capturing cannot occur when applying a β -reduction step. In particular, that means that one can follow the first, i.e., the simplest approach of the three discussed previously to define or implement substitution. In other words, if one renames the binding variables each time before applying a β -reduction, then

one does not need to rename binding variables during substitution. This is so convenient in the theoretical developments of λ -calculus, that most of the works on this subject make the following

Convention. All the binding variables occurring in any given λ -expression at any given moment are assumed to be different. Moreover, it is assumed that a variable cannot occur both free and bound in any λ -expression.

If a λ -expression does not satisfy the above convention then one can apply a certain number of α -conversions and eventually transform it into an α -equivalent one that does satisfy it.

Clearly, this process of renaming potentially all the binding variables before applying any β -reduction step may be quite expensive. In a more familiar setting, it is like traversing and changing the names of all the variables in a program at each execution step! There are techniques aiming at minimizing the

amount of work to be performed in order to avoid variable captures. All these techniques, however, incur certain overheads.

One should not get tricked by thinking that one renaming of the binding variables, at the beginning of the reduction process, should be sufficient. It is sufficient for just one step of β -reduction, but not for more. Consider, e.g., the closed λ -expression, or the combinator,

$$(\lambda z.zz)(\lambda x.\lambda y.xy).$$

It has three binding variables, all different. However, if one applies substitution in β -reductions blindly then one quickly ends up capturing the variable y :

$$\begin{array}{lcl}
\underline{(\lambda z.zz)(\lambda x.\lambda y.xy)} & \Rightarrow_{\beta} & \\
\underline{(\lambda x.\lambda y.xy)(\lambda x.\lambda y.xy)} & \Rightarrow_{\beta} & \\
\underline{\lambda y.(\lambda x.\lambda y.xy)y} & \Rightarrow_{\beta} & \\
\lambda y.\lambda y.\textcolor{red}{y}y & &
\end{array}$$

We have enough evidence by now to understand why substitution, because of the variable capture phenomenon, is considered to be such a tricky and subtle issue by many computer scientists.

We will later see an ingenious technique to transform λ -calculus into combinatory logic which, surprisingly, eliminates the need for substitutions entirely.

Confluence of β -Reduction

Consider the λ -expression $(\lambda f.(\lambda x.f x)y)g$. Note that one has two different ways to apply β -reduction on this λ -expression:

1. $(\lambda f.(\lambda x.f x)y)g \Rightarrow_{\beta} (\lambda f.f y)g$, and
2. $(\lambda f.(\lambda x.f x)y)g \Rightarrow_{\beta} (\lambda x.g x)y$.

Nevertheless, both the intermediate λ -expressions above can be further reduced to gy by applying β -reduction.

This brings us to one of the most notorious results in λ -calculus (\Rightarrow_{β}^* is the reflexive and transitive closure of \Rightarrow_{β}):

Theorem. \Rightarrow_{β} is *confluent*. That means that for any λ -expression E , if $E \Rightarrow_{\beta}^* E_1$ and $E \Rightarrow_{\beta}^* E_2$ then there is some λ -expression E' such that $E_1 \Rightarrow_{\beta}^* E'$ and $E_2 \Rightarrow_{\beta}^* E'$. All this is, of course, *modulo α -conversion*.

The confluence theorem above says that it essentially does not matter how the β -reductions are applied on a given λ -expression. A λ -expression is called a *β -normal form* if no β -reduction can be applied on it. A λ -expression E is said to *admit a β -normal form* if and only if there is some β -normal form E' such that $E \Rightarrow_{\beta}^* E'$. The confluence theorem implies that if a λ -expression admits a β -normal form then that β -normal form is *unique modulo α -conversion*.

Note, however, that there are λ -expressions which admit no β -normal form. Consider, for example, the λ -expression $(\lambda x.xx)(\lambda x.xx)$, say *omega*, known also as the *divergent combinator*. It is easy to see that *omega* \Rightarrow_{β} *omega* and that's the only β -reduction that can apply on *omega*, so it has no β -normal form.

λ -Calculus as a Programming Language

We have seen how several programming language constructs translate naturally into λ -calculus. Then a natural question arise: can we use λ -calculus as a programming language?

The answer is yes, we can, but we first need to understand how several important programming language features can be systematically captured by λ -calculus, including functions with multiple arguments, booleans, numbers, and recursion.

Currying

Recall from mathematics that there is a bijection between $[A \times B \rightarrow C]$ and $[A \rightarrow [B \rightarrow C]]$, where $[X \rightarrow Y]$ represents the set of functions $X \rightarrow Y$. Indeed, any function $f : A \times B \rightarrow C$ can be regarded as a function $g : A \rightarrow [B \rightarrow C]$, where for any $a \in A$, $g(a)$ is defined as the function $h_a : B \rightarrow C$ with $h_a(b) = c$ if and only if $f(a, b) = c$. Similarly, any function $g : A \rightarrow [B \rightarrow C]$ can be regarded as a function $f : A \times B \rightarrow C$, where $f(a, b) = g(a)(b)$.

This observation led to the important concept called *currying*, which allows us to eliminate functions with multiple arguments from the core of a language, replacing them systematically by functions admitting only one argument as above. Thus, we say that

functions with multiple arguments are just *syntactic sugar*.

From now on we may write λ -expressions of the form $\lambda xyz \dots .E$ as shorthands for their *uncurried versions* $\lambda x.\lambda y.\lambda z.\dots .E$. With this convention, λ -calculus therefore admits multiple-argument λ -abstractions. Note, however, that unlike in many familiar languages, curried functions can be applied on fewer arguments. For example, $(\lambda xyz.E)E' \beta$ -reduces to $\lambda yz.(E[x \leftarrow E'])$. Also, since λ -application was defined to be left-associative, $(\lambda xyz.E)E_1E_2 \beta$ -reduces to $\lambda z.((E[x \leftarrow E_1])[y \leftarrow E_2])$.

Most functional languages today support curried functions. The advantage of currying is that one only needs to focus on defining the meaning or on implementing effectively functions of one argument. A syntactic desugaring transformer can apply uncurrying automatically before anything else is defined.

Church Booleans

Booleans are perhaps the simplest data-type that one would like to have in a programming language. λ -calculus so far provides no explicit support for booleans or conditionals. We next show that λ -calculus provides *implicit* support for booleans. In other words, the machinery of λ -calculus is powerful enough to simulate booleans and what one would normally want to do with them in a programming language. What we discuss next is therefore a *methodology* to program with “booleans” in λ -calculus.

The idea is to regard a boolean through a “behavioral” prism: with a boolean, one can always choose one of any two objects – if true then the first, if false then the second. In other words, one can identify a boolean b with a universally quantified conditional “for any x and y , *if b then x else y* ”. With this behavior of

booleans in mind, one can now relatively easily translate booleans and boolean operations in λ -calculus:

$\text{true} := \lambda xy.x$

$\text{false} := \lambda xy.y$

$\text{if-then-else} := \lambda xyz.xyz$

$\text{and} := \lambda xy.(x\ y\ \text{false})$

This encoding for booleans is known under the name of *Church booleans*. One can use β -reduction to show, for example, that $\text{and}\ \text{true}\ \text{false} \Rightarrow_{\beta} \text{false}$. Therefore, $\text{and}\ \text{true}\ \text{false} \equiv_{\beta} \text{false}$.

One can show relatively easily that the Church booleans have all the desired properties of booleans, for example:

$\text{and}\ (\text{and}\ x\ y)\ z \equiv_{\beta} x\ y\ \text{false}\ z\ \text{false}$

$\text{and}\ x\ (\text{and}\ y\ z) \equiv_{\beta} x\ (y\ z\ \text{false})\ \text{false}$

Pairs

λ -calculus can also naturally encode data-structures of interest in most programming languages. The idea is that λ -abstractions, by their structure, can store useful information. Let us, for example, consider pairs as special cases of “records”.

Like booleans, pairs can also be regarded behaviorally: a pair is a “black-box” that can store any two expressions and then allow one to retrieve those through appropriate projections.

Formally, we would like to define λ -expressions `pair`, `1st` and `2nd` in such a way that for any other λ -expressions x and y , it is the case that `1st (pair x y)` and `2nd (pair x y)` are β -equivalent

to x and y , respectively.

Fortunately, these can be defined quite easily:

```
pair :=  $\lambda xyb.bxy$ ,  
1st  :=  $\lambda p.p \text{ true}$ , and  
2nd  :=  $\lambda p.p \text{ false}$ .
```

The idea is therefore that `pair x y` gets evaluated to the λ -expression $\lambda b.bxy$, which “freezes” x and y inside a λ -abstraction, together with a handle, b , which is expected to be a Church boolean, to “unfreeze” them later. Indeed, the first projection, `1st`, takes a pair and applies it to `true` hereby “unfreezing” its first component, while the second projection applies it to `false` to “unfreeze” its second component.

Church Numerals

Numbers and the usual operations on them can also be defined as λ -expressions. The basic idea is to regard a natural number n as a λ -expression that has the potential to apply a given operation n times on a given starting λ -expression. Therefore, λ -numerals, also called *Church numerals*, take two arguments, “what to do” and “what to start with”, and apply the first as many times as the intended numeral on the second. Intuitively, if the action was “successor” and the starting expression was “zero”, then one would get the usual numerals. Formally, we define numerals as follows:

$$0_\lambda := \lambda sz.z$$

$$1_\lambda := \lambda sz.sz$$

$$2_\lambda := \lambda sz.s(sz)$$

$$3_\lambda := \lambda sz.s(s(sz)) \quad \dots$$

With this intuition for numerals in mind, one can now easily define a successor operation on numerals:

$$\text{succ} := \lambda nsz. ns(sz)$$

The above says that for a given numeral n , its successor “ $\text{succ } n$ ” is the numeral that applies the operation s for n times starting with sz .

One can also define addition as a λ -abstraction, e.g., as follows:

$$\text{plus} := \lambda mnsz. ms(nsz)$$

One of the most natural questions that one can and should ask when one is exposed to a new model of natural numbers, is whether it satisfies the Peano axioms. In our case, this translates to whether the following properties hold:

$$\text{plus } 0_\lambda \ m_\lambda \equiv_\beta m_\lambda \text{ , and}$$

$$\text{plus } (\text{succ } n_\lambda) \ m_\lambda \equiv_\beta \text{succ } (\text{plus } n_\lambda \ m_\lambda) .$$

Interestingly, Church numerals in combination with pairs allow us to define certain recursive behaviors. Let us next define a more interesting function on Church numerals, namely one that calculates Fibonacci numbers. More precisely, we want to define a λ -expression `fibonacci` with the property that `fibonacci n_λ` β -reduces to the n -th Fibonacci number. The trick is to define a two-number “window” that “slides” through the sequence of Fibonacci numbers until it “reaches” the desired number. The window is defined as a pair and the sliding by moving the second element in the pair on the first position and placing the next Fibonacci number as the second element. The shifting operation needs to be applied as many times as the index of the desired Fibonacci number:

```
start := pair 0 $\lambda$  1 $\lambda$ ,
step  :=  $\lambda p$  . pair (2nd  $p$ ) (plus (1st  $p$ ) (2nd  $p$ )),
fibonacci :=  $\lambda n$  . 1st ( $n$  step start).
```

Another interesting use of the technique above is in defining the predecessor operation on Church numerals:

```
start := pair 0λ 0λ,
step  := λp . pair (2nd p) (plus 1λ (2nd p)),
pred  := λn . 1st (n step start).
```

Note that $\text{pred } 0_\lambda \equiv_\beta 0_\lambda$, which is a slight violation of the usual properties of the predecessor operation on integers.

The above definition of predecessor is computationally very inefficient. Unfortunately, there does not seem to be any better way to define this operation on Church numerals.

Subtraction can now be defined easily:

```
sub := λmn. n pred m.
```

Note, again, that negative numbers are collapsed to 0_λ .

Let us next see how relational operators can be defined on Church numerals. These are useful to write many meaningful programs. We first define a helping operation, to test whether a number is zero:

`zero? := λn . n (and false) true.`

Now the “less than or equal to” (`leq`), the “larger than or equal to” (`geq`), and the “equal to” (`equal`) can be defined as follows:

`leq := λmn . zero? (sub m n),`

`geq := λmn . zero? (sub n m),`

`equal := λmn . and (leq m n) (geq m n).`

η-Equivalence

Let us consider the λ -expression $\lambda x.Ex$, where E is some λ -expression that does not contain x free. Intuitively, $\lambda x.Ex$ does nothing but wraps E : when “called”, it “passes” its argument to E and then “passes” back E ’s result. When applied on some λ -expression, say E' , note that $\lambda x.Ex$ and E *behave the same*. Indeed, since E does not contain any free occurrence of x , one can show that $(\lambda x.Ex)E' \equiv_{\beta} EE'$. Moreover, if E is a λ -abstraction, say $\lambda y.F$, then $\lambda x.Ex = \lambda x.(\lambda y.F)x \equiv_{\beta} \lambda x.F[y \leftarrow x]$. The latter is α -equivalent to $\lambda y.F$, so it follows that in this case $\lambda x.Ex$ is β -equivalent to E .

Even though $\lambda x.Ex$ and E have similar behaviors in applicational contexts and they can even be shown β -equivalent when E is a λ -abstraction, there is nothing to allow us to use their equality as

an axiom in our equational inferences. In particular, there is no way to show that the combinator $\lambda x.\lambda y.xy$ is equivalent to $\lambda x.x$.

To increase the proving capability of λ -calculus, still without jeopardizing its basic intuitions and applications, we consider its extension with the following equation:

$$(\eta) \qquad \lambda x.Ex = E,$$

for any $x \notin FV(E)$. We let $E \equiv_{\beta\eta} E'$ denote the fact that the equation $E = E'$ can be proved using standard equational deduction from all the equational axioms above: (α) , (β) , (η) , plus those for substitution. The relation $\equiv_{\beta\eta}$ is also called *$\beta\eta$ -equivalence*. The λ -calculus enriched with the rule (η) is also called $\lambda + \eta$.

Extensionality

Extensionality is a deduction rule encountered in several branches of mathematics and computer science. It intuitively says that in order to prove two objects equal, one may first “extend” them in some rigorous way. The effectiveness of extensionality comes from the fact that it may often be the case that the extended versions of the two objects are easier to prove equivalent.

Extensionality was probably first considered as a proof principle in set theory. In “naive” set theory, sets are built in a similar fashion to Peano numbers, that is, using some simple constructors (together with several constraints), such as the empty set \emptyset and the list constructor $\{x_1, \dots, x_n\}$. Thus, $\{\emptyset, \{\emptyset, \{\emptyset\}\}\}$ is a well-formed set.

With this way of constructing sets, there may be the case that two

sets with “the same elements” have totally different representations. Consequently, it is almost impossible to prove any meaningful property on sets, such as distributivity of union and intersection, etc., by just taking into account how sets are constructed. In particular, proofs by structural induction are close to useless.

Extensionality is often listed as the first axiom in any axiomatization of set theory. In that context, it basically says that two sets are equal iff they have the same elements. Formally,

If $x \in S = x \in S'$ for any x , then $S = S'$.

Therefore, in order to show sets S and S' equal, one can first “extend” them (regarded as syntactic terms) by applying them the membership operator. In most cases the new task is easier to prove.

In λ -calculus, extensionality takes the following shape:

(ext) If $Ex = E'x$ for some $x \notin FV(EE')$, then $E = E'$.

Therefore, two λ -abstractions are equal if they are equal when applied on some variable that does not occur free in any of them.

Note that “for some x ” can be replaced by “for any x ” in **ext**. We let $E \equiv_{\beta\text{ext}} E'$ denote the fact that the equation $E = E'$ can be proved using standard equational deduction using (α) and (β) , together with **ext**. The λ -calculus extended with **ext** is also called $\lambda + \text{ext}$.

The following important result says the extensions of λ -calculus with (η) and with **ext** are equivalent:

Theorem. $\lambda + \eta$ is equivalent to $\lambda + \text{ext}$.

Combinatory Logic

Even though λ -calculus can be defined equationally and is a relatively intuitive framework, as we have noticed several times by now, substitution makes it non-trivial to implement effectively. There are several approaches in the literature addressing the subtle problem of automating substitution to avoid variable capture, all with their advantages and disadvantages. We here take a different approach. We show how λ -expressions can be automatically translated into expressions over combinators, in such a way that substitution will not even be needed anymore.

A question addressed by many researchers several decades ago, still interesting today and investigated by many, is whether there is any *simple* equational theory that is entirely equivalent to λ -calculus. Since λ -calculus is Turing complete, such a simple theory may

provide a strong foundation for computing.

Combinatory logic was invented by Moses Shönfinkel in 1920. The work was published in 1924 in a paper entitled “*On the building blocks of mathematical logic*”. Combinatory logic is a simple equational theory over two sorts, *Var* and *Exp* with $\text{Var} < \text{Exp}$, a potentially infinite set x, y , etc., of constants of sort *Var* written using lower-case letters, two constants K and S of sort *Exp*, one application operation with the same syntax and left-associativity parsing convention as in λ -calculus, together with the two equations

$$\begin{aligned} KXY &= X, \\ SXYZ &= XZ(YZ), \end{aligned}$$

quantified universally over X, Y, Z of sort *Exp*. The constants K and S are defined equationally in such a way to capture the intuition that they denote the combinators $\lambda xy.x$ and $\lambda xyz.xz(yz)$, respectively. The terms of the language, each of which denoting a

function, are formed from variables and constants K and S by a single construction, function application. For example, $S(SxKS)yS(SKxK)z$ is a well-formed term in combinatory logic, denoting some function of free variables x , y , and z .

Let \mathbf{CL} be the equational theory of combinatory logic above. Note that a function FV returning the “free” variables that occur in a term in combinatory logic can be defined in a trivial manner, because there are no “bound” variables in \mathbf{CL} . Also, note that the extensionality principle from λ -calculus translates unchanged to \mathbf{CL} :

(ext) If $Ex = E'x$ for some $x \notin FV(EE')$, then $E = E'$.

Let $\mathbf{CL} + \mathbf{ext}$ be \mathbf{CL} enriched with the principle of extensionality. The following is a landmark result:

Theorem. $\lambda + \mathbf{ext}$ is equivalent to $\mathbf{CL} + \mathbf{ext}$.

Proof. Let us recall what one needs to show in order for two

mathematical theories to be equivalent: (1) how the syntax of one translates into the syntax of the other; and (2) that all the axioms of each of the two theories can be proved from the axioms of the other, along the corresponding translation of syntax.

Let us consider first the easy part: $\lambda + \text{ext}$ implies $\text{CL} + \text{ext}$. We first need to show how the syntax of $\text{CL} + \text{ext}$ translates into that of $\lambda + \text{ext}$. This is easy and it was already mentioned before: let K be the combinator $\lambda xy.x$ and let S be the combinator $\lambda xyz.xz(yz)$. We then need to show that the two equational axioms of $\text{CL} + \text{ext}$ hold under this translation: they can be immediately proved by β -equivalence. We also need to show that the extensionality in $\text{CL} + \text{ext}$ holds under the above translation: this is obvious, because it is exactly the same as the extensionality in $\lambda + \text{ext}$.

Let us now consider the other, more difficult, implication. So we start with **CL + ext**, where K and S have no particular meaning in λ -calculus, and we need to define some map that takes any λ -expression and translates it into an expression in **CL**.

To perform such a transformation, let us add syntax for λ -abstractions to **CL**, but without any of the equations of λ -calculus. This way one can write and parse λ -expressions, but still have no meaning for those. The following ingenious rewriting system transforms any uninterpreted λ -expression into an expression using only K , S , and the free variables of the original λ -expression:

$$1. \lambda x.y \Rightarrow \begin{cases} SKK & \text{if } x = y \\ Ky & \text{if } x \neq y \end{cases}$$

$$2. \lambda x.(\rho\rho') \Rightarrow S(\lambda x.\rho)(\lambda x.\rho')$$

$$3. \lambda x.K \Rightarrow KK$$

$$4. \lambda x.S \Rightarrow KS$$

Here ρ and ρ' can be any expressions over K , S , variables, and the application operator, but also over the λ -abstraction operator $\lambda_{_} : Var \rightarrow Exp$. The rules above systematically eliminate all the brackets. Therefore, the “bracket abstraction” rules above eventually transform any λ -expression into an expression over only K , S , variables, and the application operator.

The correctness of the translation of $\lambda + \text{ext}$ into $\text{CL} + \text{ext}$ via the technique above is rather technical: one needs to show that the translated versions of equations in λ can be proved (by structural induction) using the machinery of $\text{CL} + \text{ext}$. We do not need to understand the details of the proof of correctness in order to have a good intuition on why the bracket abstraction translation works. To see that, just think of the translation technique above as a means to associate equivalent λ -expressions to other λ -abstractions, within the framework of λ -calculus, where K and S are their

corresponding λ -expressions. As seen above, it eventually reduces any λ -expression to one over only combinators and variables, containing no explicit λ -abstractions except those that define the combinators K and S . To see that the translation is correct, we can show that each rule is sound within λ -calculus. For example, rule 2 can be shown by extensionality:

$$(\lambda x.\rho\rho')z \equiv_{\beta} (\rho[x \leftarrow z])(\rho'[x \leftarrow z]) \equiv_{\beta} ((\lambda x.\rho)z)((\lambda x.\rho')z) \equiv_{\beta} (\lambda xyz.xz(yz))(\lambda x.\rho)(\lambda x.\rho')z = S(\lambda x.\rho)(\lambda x.\rho')z,$$

so by extensionality, $\lambda x.\rho\rho' \equiv_{\beta_{\text{ext}}} S(\lambda x.\rho)(\lambda x.\rho')$.

This way, one can prove the soundness of each of the rules in the bracket abstraction translation. As one may expect, the tricky part is to show the completeness of the translation, that is, that everything one can do with λ -calculus and **ext** can also do with its “sub-calculus” **CL** + **ext**. This is not hard, but rather technical.